UNITED STATES PATENT AND TRADEMARK OFFICE

A

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO:	CONFIRMATION NO.
10/780,772	02/19/2004	Kentaro Shimada	H-1131	2065

24956 7590 09/21/2005

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.
1800 DIAGONAL ROAD
SUITE 370
ALEXANDRIA, VA 22314

EXAMINER	
FLOURNOY, HORACE L	

ART UNIT	PAPER NUMBER
2189	

DATE MAILED: 09/21/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/780,772	SHIMADA ET AL.	
	Examiner	Art Unit	
	Horace L. Flournoy	2189	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *19 February 2004*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-18* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-18* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☒ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on *19 February 2004* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☒ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☒ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☒ Certified copies of the priority documents have been received in Application No. 10/780,772.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other: _____.

DETAILED ACTION

Claims 1-18 are presented for examination.

Priority

Acknowledgment is made of applicant's claim for foreign priority under 35 U.S.C. 119(a)-(d). The certified copy has been filed in parent Application No. 2003-418907, filed on 12/17/2003.

Claim Rejections - 35 USC § 112

The following is a quotation of the first paragraph of 35 U.S.C. 112:

> The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

Claims 9 recites "attaches importance." The specification does not enable one of ordinary skill in the art to know how to "attach importance" to something. In other words, what is "important?" What decides whether something is important? How is that importance "attached?"

The following is a quotation of the second paragraph of 35 U.S.C. 112:

> The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

Claims 9, 13-15 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

Claims 9 recites "attaches importance." The specification does not particularly point out what this means.

Claim 13 recites "accesses a small number of large files." Claim 14 recites "accesses a large number of small files." Claim 15 recites "accesses a large file." In each of these limitations the meaning of "large" or "small" is relative to an unknown value.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 1-18 are rejected under 35 U.S.C. 102(e) as being anticipated by Doing et al. (U.S. PG Pub no. 2003/0009648, hereafter referred to as Doing).

With respect to independent claims 1, 16, and 18,

"A storage to be connected to a network, comprising:

a plurality of interfaces which is connected to the network and receives file access;

a plurality of disk drives;" is disclosed on page 1, paragraph [0018].

The examiner interprets the limitation "...receives files access..." to mean storing, retrieving, and transferring (access) information (file).

Doing discloses on page 1, paragraph [0018], "A modern computer system typically comprises a central processing unit (CPU) and supporting hardware necessary to store, retrieve and transfer information, such as communications busses and memory. It also includes hardware necessary to communicate with the outside world, such as input/output controller or storage controllers and devices attached thereto such as keyboards, monitors, tape drives, disk drives, communication lines coupled to a network, etc." Doing teaches multiple storage devices connected to a network including disk drives and tape drives. Doing also teaches a plurality of interfaces which is connected to the network and receives file access (hardware necessary to communicate with the outside world, such as input/output controller or storage controllers).

"...and a control unit..." is disclosed on page 3, paragraph [0035].

The examiner interprets control unit (based of off its functional description and claim limitations) to be a hypervisor, partition manager, or a (set of) device(s) that controls the logical partitioning of a computer system.

Doing discloses in paragraph [0035] "A processor provides hardware support for logical partitioning of a computer system. Logical partitions isolate the real address spaces of processes executing on different processors, specifically, supervisory processes. An ultra-privileged supervisor process, called a hypervisor, regulates the logical partitions."

"...which translates the file access into block access..." is disclosed in paragraphs [0037] and [0047].

The limitation "...translates the file access into block access..." is interpreted as translating effective or virtual addresses into real or physical addresses.

Doing next discloses in paragraph [0037], "...each processor generates effective addresses from executable code, which are translated to real addresses corresponding to locations in physical memory." Doing discloses in paragraph [0047] "FIG. 4 illustrates in greater detail real address partitioning logic, an effective to real address table and associated control structures for instruction addresses..."

"...and controls the plurality of disk drives on the basis of the block access..." is disclosed in paragraph [0035].

The limitation "...controls the plurality of disk drives on the basis of the block access..." is interpreted as controlling the plurality of disk drives (which are part of the plurality of virtual partitions) on the basis of real or physical addresses.

Doing discloses in paragraph [0035] "A processor provides hardware support for logical partitioning of a computer system. Logical partitions isolate the real address spaces of processes executing on different processors, specifically, supervisory processes. An ultra-privileged supervisor process, called a hypervisor, regulates the logical partitions." The plurality of disk drives (logical partitions) is controlled (hardware support for logical partitioning) on the basis of block access (logical partitions isolate the real address spaces of processes executing on different processors).

"...wherein the control unit logically partitions the plurality of interfaces, the plurality of disk drives and the control unit and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives and the partitioned control unit to operate as a plurality of virtual storages independently." is disclosed on page 3, paragraph [0029].

The examiner interprets control unit (based of off its functional description and claim limitations) to be a hypervisor, partition manager, or a (set of) device(s) that controls the logical partitioning of a computer system.

Doing discloses on page 3, paragraph [0029], "Hardware support would also be useful for dynamically re-partitioning the system in an efficient manner. This hardware support may be used to enforce the partitioning of system resources such as

processors, real memory, internal registers, etc." Doing teaches the control unit (hardware support) that logically partitions the plurality of interfaces (system resources), the partitioned plurality of disk drives (real memory) and the control unit (processors). Doing further teaches the independent virtual storages or partitions in paragraph [0137]: "Each logical partition has its own segment and page tables, which are maintained independently of those in other logical partitions."

With respect to claim 2, and independent claim 18,

> "A storage according to claim 1, wherein the control unit further includes a plurality of cache memories, and the plurality of cache memories is logically partitioned and allocated to the respective plurality of virtual storages" is disclosed in paragraphs [0053] and [0065] and FIG. 1 and 2.

The examiner interprets control unit (based of off its functional description and claim limitations) to be a hypervisor, partition manager, or a (set of) device(s) that controls the logical partitioning of a computer system. The examiner interprets the limitation, "...and the plurality of cache memories is logically partitioned and allocated to the respective plurality of virtual storages" to mean the cache memories (of multiple processors which constitute the control unit) are logically partitioned and allocated to the respective plurality of virtual storages.

Doing discloses in paragraph [0065], "In the preferred embodiment, processors and real memory are assigned to logical partitions in a partitioned system." In

paragraph [0053], Doing discloses, "CPUs 101A, 101B, 101C and 101D for processing instructions contains separate respective internal level one instruction caches 106A, 106B, 106C, 106D (L11-cache) and level one data caches 107A, 107B, 107C, 107D (L1D-cache). Each L11-cache 106A, 106B, 106C, 106D stores instructions for execution by its CPU 101A, 101B, 101C, 101D. L I D-cache stores data (other than instructions) to be processed by a CPU. Each CPU 101A, 101B, 101C, 101D is coupled to a respective level two cache (L2 cache) 108A, 108B, 108C, 108D, which can be used to hold both instructions and data." FIG. 1 teaches the implementation of the control unit (multiple processors) and FIG. 2 teaches the implementation of a plurality of cache memories. Since Doing is partitioning the plurality of processors, the plurality of cache memories are also logically partitioned and allocated to the respective plurality virtual storages.

With respect to claim 3, and independent claim 18,

> "A storage according to claim 2, wherein the control unit further includes a first processor which translates the file access into the block access..." is disclosed in paragraphs [0037], [0047], and [0060].

The limitation "...translates the file access into block access..." is interpreted as translating effective or virtual addresses into real or physical addresses.

Doing discloses in paragraph [0060], "CPU 101 of the preferred embodiment supports multiple levels of address translation, as logically illustrated in FIG. 8. The

three basic addressing constructs are effective address 801, virtual address 802, and real address 803." Doing next discloses in paragraph [0037], "...each processor generates effective addresses from executable code, which are translated to real addresses corresponding to locations in physical memory." Doing discloses in paragraph [0047] "FIG. 4 illustrates in greater detail real address partitioning logic, an effective to real address table and associated control structures for instruction addresses..."

"...and a second processor, which controls the plurality of disk drives on the basis of the block access..." is disclosed in paragraphs [0035], [0036], and [0065].

The examiner interprets this limitation to mean the second (or separate) processor controls plurality of disk drives (disk drives of the virtual partitions) based on or using real addresses

Doing discloses in paragraph [0035] "A processor provides hardware support for logical partitioning of a computer system. Logical partitions isolate the real address spaces of processes executing on different processors, specifically, supervisory processes. An ultra-privileged supervisor process, called a hypervisor, regulates the logical partitions." In paragraph [0036] Doing teaches that "...the processor contains multiple register sets for supporting the concurrent execution of multiple threads (i.e., hardware multithreading). Each thread is capable of independently being in either hypervisor, supervisor or problem (non-privileged) state." Doing teaches in paragraph

[0065] that "...the different logical partitions share hardware resources such as disk storage...In a logically partitioned system, each processor of the multiprocessor system is assigned to a partition, along with a subset of the real memory address space." Each processor has the ability to function as a hypervisor, which regulates the logical partitions (controls the plurality of disk drives) on the basis of block access (logical partitions isolate the real address spaces of processes executing on different processors).

> "...and wherein the first processor and the second processor are logically partitioned respectively, and allocated to the respective plurality of virtual storages" is disclosed in paragraph [0065].

The examiner interprets this limitation to mean, the first processor and the second processor (different processors) are logically partitioned respectively, and allocated to the respective plurality of virtual storages (virtual partitions).

Doing discloses in paragraph [0065], "In the preferred embodiment, processors and real memory are assigned to logical partitions in a partitioned system...In a logically partitioned system, each processor of the multiprocessor system is assigned to a partition, along with a subset of the real memory address space."

With respect to claims 4, 6, and 7,

"A storage according to claim 3, wherein the first processor executes first hypervisor which performs logical partitioning of the plurality of interfaces and the first processor..." is disclosed in paragraphs [0035] and [0036].

The examiner interprets this claim to mean the first processor executes a hypervisor which performs logical partitioning of the plurality of interfaces and the first processor.

In paragraph [0036] Doing teaches that "...the processor contains multiple register sets for supporting the concurrent execution of multiple threads (i.e., hardware multithreading). Each thread is capable of independently being in either hypervisor, supervisor or problem (non-privileged) state." Doing teaches that each processor is capable of executing a hypervisor.

Doing discloses in paragraph [0035], "An ultra-privileged supervisor process, called a hypervisor, regulates the logical partitions."

In paragraph [0036] Doing teaches that "...the processor contains multiple register sets for supporting the concurrent execution of multiple threads (i.e., hardware multithreading). Each thread is capable of independently being in either hypervisor, supervisor or problem (non-privileged) state." Since each processor supports multi-threading, each thread within a particular processor can be logically partitioned. As stated supra, logical partitioning is handled by the hypervisor.

"...and wherein the second processor executes second hypervisor which performs logical partitioning of the plurality of cache memories, the plurality of

disk devices and the second processor" is disclosed in paragraphs [0029] and [0036].

The examiner interprets this claim to mean the second processor executes another hypervisor which performs logical partitioning of the plurality of cache memories, the plurality of disk devices and the second processor.

In paragraph [0036] Doing teaches that "...the processor contains multiple register sets for supporting the concurrent execution of multiple threads (i.e., hardware multithreading). Each thread is capable of independently being in either hypervisor, supervisor or problem (non-privileged) state." Doing teaches that each processor is capable of executing a hypervisor.

FIG. 2 teaches that the plurality of cache memories (L1 and L2) is divided by respective processors. Doing discloses on page 3, paragraph [0029], "Hardware support would also be useful for dynamically re-partitioning the system in an efficient manner. This hardware support may be used to enforce the partitioning of system resources such as processors, real memory, internal registers, etc."

Doing discloses that hardware support can dynamically re-partition the system resources including the processors, real memory, internal registers, etc. The plurality of cache memories is therefore logically partitioned since they are a part of the plurality of processors. For similar reasons, plurality of disk devices (real memory) is also logically partitioned by a hypervisor (which a second processor can execute).

In paragraph [0036] Doing teaches that "...the processor contains multiple register sets for supporting the concurrent execution of multiple threads (i.e., hardware

multithreading). Each thread is capable of independently being in either hypervisor, supervisor or problem (non-privileged) state." Since each processor supports multi-threading, each thread within a particular processor can be logically partitioned. As stated supra, logical partitioning is handled by the hypervisor.

With respect to claim 5, and independent claim 18,

> "A storage according to claim 4, wherein the control unit further includes a plurality of memories which is used by the first processor and a plurality of communication units which connects the first processor and the second processor..." is disclosed in FIG. 1 and paragraph [0053].

The examiner interprets this claim to mean the control unit (multiple processors) further includes a plurality of memories (106A and 108A, L1 and L2 caches for of FIG. 1) which is used by the first processor (CPU 101A) and a plurality of communication units which connects the first processor and the second processor.

In paragraph [0053] Doing teaches that "CPUs 101A, 101B...communicate via bus interface 105 with system bus 110." FIG. 1 shows the plurality of communication buses (elements 109 and 105) between the first (CPU 101A) and second (CPU 101B) processors.

> "...wherein the plurality of memories is logically partitioned by the first hypervisor..." is disclosed as stated supra.

"...and the plurality of communication units is logically partitioned by the second hypervisor" is disclosed in paragraphs [0035] and [0036].

The examiner interprets this claim to mean the plurality of communication units is logically partitioned by the second hypervisor.

Doing teaches in paragraph [0053], "Various I/O processing units (IOPs) 111-115 attach to system bus 110 and support communication..." The hypervisor (which can be executed by any of the multi-threaded processors in the control unit) controls logical partitioning of the system resources (IOPs).

With respect to claims 8, 11, and 16,

"A storage according to claim 3 further connected to a supervising terminal, wherein the control unit performs the logical partitioning on the basis of information inputted from the supervising terminal" is disclosed in paragraph [0018].

The examiner interprets supervising terminal to mean a (set of) hardware devices that provide communication to the NAS.

In paragraph [0018] Doing discloses that "A modern computer system typically comprises a central processing unit (CPU) and supporting hardware necessary to store, retrieve and transfer information....It also includes hardware necessary to communicate with the outside world, such as input/output controllers or storage controllers, and devices attached thereto such as keyboards, monitors...The CPU is the heart of the

system. It executes the instructions which comprise a computer program and directs the operation of the other system components."

Doing teaches that the system which includes the CPU and supporting hardware (NAS) also include hardware such as keyboards and monitors that are *necessary* for communicating with the outside world. The supervising terminal (keyboard, monitor, and other I/O hardware) is where information for performing logical partitioning is originated, since the NAS system executes the instruction from a computer program.

With respect to claims 9 and 12,

> "A storage according to claim 8, wherein, if information to be inputted to the supervising terminal is information to the effect that a host system using the storage attaches importance to a data transfer rate, an amount of allocation of the plurality of cache memories to a virtual storage to be used by the host system among the plural virtual storages is increased" is disclosed in paragraph [0084], FIGs 1-3, and 7.

The examiner interprets claim 9 to mean if the instruction (from a supervising terminal) has an importance in data transfer rate, then increase the amount of allocation of shared cache to a virtual partition. The examiner interprets the limitation "attaches importance" to mean an instruction that depends mainly on or is performance enhanced by data transfer rate.

Doing discloses in paragraph [0084] "Ideally, instruction unit 201 provides a constant stream of instructions for decoding in decoder 206, and execution by execution unit 211. L11-cache 106 must respond to an access request with minimal delay. Where a requested instruction is actually in L11-cache, it must be possible to respond and fill the appropriate buffer without requiring decoder/dispatcher 206 to wait. Where L11-cache cannot respond (i.e., the requested instruction is not in L11-cache), a longer path via cache fill bus 233 through memory management unit 222 must be taken. In this case, the instruction may be obtained from L2 cache 108, from main memory 102, or potentially from disk or other storage. It is also possible that the instruction will be obtained from L2 cache of another processor. In all of these cases, the delay required to fetch the instruction from a remote location may cause instruction unit 201 to switch threads. I.e., the active thread becomes inactive, the previously inactive thread becomes active, and the instruction unit 201 begins processing instructions of the previously inactive thread held in thread switch buffer 204."

Doing teaches, the instruction (from a supervising terminal) has an importance for data transfer rate(Where L11-cache cannot respond (i.e., the requested instruction is not in L11-cache), a longer path via cache fill bus 233 through memory management unit 222 must be taken.), then increase the amount of allocation of shared cache to a virtual partition (It is also possible that the instruction will be obtained from L2 cache of another processor). Doing shows that the amount of allocation of the plurality of cache memories (L1/L2 cache within each processor) can increase (a processor can use another's cache in addition to its own).

With regard to claim 12, Doing also teaches the limitation "sequential continuous access" (constant stream of instructions). Doing further teaches "an amount of allocation of the plurality of cache memories and the plurality of memories which is used by the first processor" (the instruction may be obtained from L2 cache 108, from main memory 102, or potentially from disk or other storage).

With respect to claim 10,

> "A storage according to claim 8, wherein, if information to be inputted to the supervising terminal is information to the effect that a host system using the storage performs random access in a large area, an amount of allocation of the plurality of cache memories to a virtual storage to be used by the host system among the plural virtual storages is reduced" is disclosed in paragraph [0084], FIGs 1-3, and 7.

The examiner interprets claim 10 to mean if the instruction (from a supervising terminal) has an importance in response time, then reduce the amount of allocation of shared cache to a virtual partition. The examiner interprets the limitation, "performs random access in a large area" to mean response time based on the specification of this application.

Doing discloses in paragraph [0084] "Ideally, instruction unit 201 provides a constant stream of instructions for decoding in decoder 206, and execution by execution unit 211. L11-cache 106 must respond to an access request with minimal delay.

Where a requested instruction is actually in L11-cache, it must be possible to respond and fill the appropriate buffer without requiring decoder/dispatcher 206 to wait. Where L11-cache cannot respond (i.e., the requested instruction is not in L11-cache), a longer path via cache fill bus 233 through memory management unit 222 must be taken. In this case, the instruction may be obtained from L2 cache 108, from main memory 102, or potentially from disk or other storage. It is also possible that the instruction will be obtained from L2 cache of another processor. In all of these cases, the delay required to fetch the instruction from a remote location may cause instruction unit 201 to switch threads. I.e., the active thread becomes inactive, the previously inactive thread becomes active, and the instruction unit 201 begins processing instructions of the previously inactive thread held in thread switch buffer 204."

Doing teaches, the instruction (from a supervising terminal) has an importance in response time (Where L11-cache cannot respond (i.e., the requested instruction is not in L11-cache), a longer path via cache fill bus 233 through memory management unit 222 must be taken), then reduce the amount of allocation of shared cache to a virtual partition. Doing teaches that an allocation of cache is reduced for a particular partition (processor) when another processor is utilizing it.

With respect to claims 13 and 14,

"A storage according to claim 8, wherein, if information to be inputted to the supervising terminal is information to the effect that a host system using the storage accesses a small number of large files, an amount of allocation of the first processor to a virtual storage to be used by the host system among the plural virtual storages is reduced, and an amount of allocation of the second processor to the virtual storage is increased.

and

"A storage according to claim 8, wherein, if information to be inputted to the supervising terminal is information to the effect that a host system using the storage accesses a large number of small files, an amount of allocation of the first processor to the a virtual storage to be used by the host system among the plural virtual storages is increased, and an amount of allocation of the second processor to the virtual storage is reduced" is disclosed in Armstrong et at., column 4, lines 35-67(U.S. Pat. no. 6,279,046, incorporated by reference by Doing and hereafter referred to as Armstrong).

The examiner interprets the limitation "accesses a small number of large files" of claim 13 and "accesses a large number of small files" of claim 14 as inherent to accessing storage.

Armstrong discloses in column 4, lines 35-46, "Each logical partition 40-44 executes in a separate, or independent, memory space, represented by virtual memory

60. Moreover, each logical partition 40-44 is statically and/or dynamically allocated a portion of the available resources in computer 10. For example, each logical partition is allocated one or more processors 12, as well as a portion of the available memory space for use in virtual memory 60. Logical partitions can share specific hardware resources such as processors, such that a given processor is utilized by more than one logical partition. In the alternative hardware resources can be allocated to only one logical partition at a time."

Armstrong teaches amount of allocation of the first processor to the a virtual storage to be used by the host system among the plural virtual storages is increased, and an amount of allocation of the second processor to the virtual storage is reduced (...each logical partition is allocated one or more processors 12, as well as a portion of the available memory space for use in virtual memory 60. Logical partitions can share specific hardware resources such as processors, such that a given processor is utilized by more than one logical partition.)

With respect to claim 15,

> "A storage according to claim 11, wherein information to be inputted to the supervising terminal is information to the effect that a host system using the storage sequentially accesses a large file, an amount of logical allocation of the plurality of communication units to a virtual storage to be used by the host system reduced among the plural virtual storages is reduced" is disclosed in

Armstrong et at., column 4, lines 47-67(U.S. Pat. no. 6,279,046, incorporated by reference by Doing and hereafter referred to as Armstrong).

The examiner interprets the limitation "sequentially accesses a large file" as inherent to accessing storage devices.

Armstrong discloses in column 4, lines 47-59, "Additional resources, e.g., mass storage, backup storage, user input, network connections, and the like, are typically allocated to one or more logical partitions in a manner well known in the art. Resources can be allocated in a number of manners, e.g., on a bus-by-bus basis, or on a resource-by-resource basis, with multiple logical partitions sharing resources on the same bus. Some resources may even be allocated to multiple logical partitions at a time. FIG. 2 illustrates, for example, three logical buses 62, 64 and 66, with a plurality of resources on bus 62, including a direct access storage device (DASD) 68, a control panel 70, a tape drive 72 and an optical disk drive 74, allocated to primary logical partition 40."

With respect to claim 17,

"A storage system according to claim 16, wherein information to be inputted to the supervising terminal is information on characteristics of accesses of a computer using the storage, and the storage calculates an amount of logical partitioning of resources provided in the storage on the basis of the information on characteristics of accesses to be inputted to the supervising terminal and

performs the logical partitioning using a result of the calculation...." is disclosed in paragraphs [0040] and [0042].

The examiner interprets this claim to mean the inputted information to the supervising terminal has characteristics of accesses of a computer using the storage. The storage (NAS or system) calculates an amount of logical partitioning of resources provided in the storage on the basis of the inputted information and performs the logical partitioning using the result of the calculation.

Doing discloses in paragraph [0040], "In the preferred embodiment, special operating system software running in hypervisor state can dynamically re-allocate resources to logical partitions. In particular, it can alter the contents of the real memory offset register and the real memory limit register which regulate the generation of partitioned real addresses; a logical partition identifier which identifies the logical partition to which a processor is assigned; and certain configuration information."

Doing teaches the inputted information to the supervising terminal (via operating system software) has characteristics of accesses of a computer using the storage (running in hypervisor state). The storage (NAS or system) calculates an amount of logical partitioning of resources provided in the storage (configuration information) on the basis of the inputted information and performs the logical partitioning using the result of the calculation (dynamically re-allocate resources to logical partitions).

Conclusion

Any inquiry concerning this communication or earlier communication from the examiner should be directed to Horace L. Flournoy whose telephone number is (571) 272-2705. The examiner can normally be reached on Monday-Friday 7:00 AM to 4:30 PM (ET).

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Don Sparks can be reached on (571) 272-4201. The fax phone numbers for the organization where this application or proceeding is assigned is (703) 746-7239

Information regarding the status of an Application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or PUBLIC PAIR. Status information for unpublished applications is available through Private Pair only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (571) 272-2100.

Horace L. Flournoy
Patent Examiner
Art unit: 2189

CHRISTIAN CHACE
PRIMARY EXAMINER
Primary Patent Examiner
Technology Center 2100